

Who are the people in my neighborhood? The ‘contextual fallacy’ of measuring individual context with census geographies

by

**Christopher S. Fowler
Pennsylvania State University**

**Nathan Frey
Brown University**

**David C. Folch
Florida State University**

**Nicholas Nagle
University of Tennessee, Knoxville**

**Seth Spielman
University of Colorado, Boulder**

CES 18-11

February, 2018

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.census.gov/ces or contact J. David Brown, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 5K034A, 4600 Silver Hill Road, Washington, DC 20233, CES.Working.Papers@census.gov. To subscribe to the series, please click [here](#).

Abstract

Scholars deploy census-based measures of neighborhood context throughout the social sciences and epidemiology. Decades of research confirm that variation in how individuals are aggregated into geographic units to create variables that control for social, economic or political contexts can dramatically alter analyses. While most researchers are aware of the problem, they have lacked the tools to determine its magnitude in the literature and in their own projects. By using confidential access to the complete 2010 U.S. Decennial Census, we are able to construct—for all persons in the US—individual-specific contexts, which we group according to the Census-assigned block, block group, and tract. We compare these individual-specific measures to the published statistics at each scale, and we then determine the magnitude of variation in context for an individual with respect to the published measures using a simple statistic, the standard deviation of individual context (SDIC). For three key measures (percent Black, percent Hispanic, and Entropy—a measure of ethno-racial diversity), we find that block-level Census statistics frequently do not capture the actual context of individuals within them. More problematic, we uncover systematic spatial patterns in the contextual variables at all three scales. Finally, we show that within-unit variation is greater in some parts of the country than in others. We publish county-level estimates of the SDIC statistics that enable scholars to assess whether mis-specification in context variables is likely to alter analytic findings when measured at any of the three common Census units.

Keyword: context, race, census data

* The authors would like to thank James Gaboardi, and Parmanand Sinha for shared labors in the Census Restricted Data Center that helped to produce the results presented here. Research was funded by NSF Grant 1132008 supporting the Census Research Network: Spatial Science Node. We acknowledge assistance provided by the Population Research Institute at Pennsylvania State University, which is supported by an infrastructure grant by the Eunice Kennedy Shriver National Institute of Child Health and Human Development (P2CHD041025). Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

Introduction

The ‘ecological fallacy’ occurs when correlations between measures based on aggregate data do not apply to the individuals within the aggregating unit. Robinson famously demonstrated this problem using examples of race and nativity status correlated with illiteracy (1). His argument pushed many areas of social science research towards studies based on individual-level data and, according to Firebaugh, led to major shifts in approaches to social science research that continue to the present day (2, 3). An analogous problem arises, however, when scholars fail to recognize that individuals operate within social, political, geographic, and economic contexts that shape a wide range of outcomes from economic opportunity to health.¹ New multi-level methods combined with a ‘spatial turn’ in several of the social sciences mean that contextual effects are integral to a wide range of research agendas. The interdisciplinary literature on neighborhood effects (4–10) represents just one of many lines of inquiry that rely heavily on contextual variables.

This paper examines what we refer to as the *contextual fallacy*: the assumption that a contextual variable defined for a population within some unit of observation adequately reflects the experience of all individuals residing within that unit. To illustrate, we offer two examples:

Valid: A measure of the generosity of state-provided welfare benefits to characterize the effects on individuals of a social safety net. In this example, an individual’s home address conditions the form and quantity of benefit available, so the context of potential safety-net generosity only varies between households in different states. An individual living in the center of the state or in its largest city has, in principle at least, exactly the same access to benefits as an individual living just within the state border or in a rural part of the state.

Problematic: A measure of racial segregation defined for metropolitan regions to characterize the effects on individuals of local race relations. In this example, individuals within the same metropolitan area may experience a wide range of residential sorting outcomes, from segregated to diverse, depending on where they live within a single metropolitan region. In this case the aggregating unit is poorly suited to the context it is meant to represent and hides considerable variation in experience for individuals within the metropolitan area. It is even possible that the metropolitan average does not fit a single individual within the region.

¹ in reaction to Robinson and the resulting over-emphasis on individual characteristics as determinative of outcomes, the failure to acknowledge these contextual effects is sometimes referred to as the ‘individual’ or ‘atomistic’ fallacy c.f. (3, 10)

The problem identified in the preceding examples can be solved by selecting appropriate units of aggregation. However, for many social science questions the context of interest may not have definite borders, as in the case of neighborhoods (11, 12), or labor markets (13, 14).

The contextual fallacy is related to Robinson's work on the ecological fallacy, but is more accurately identified as a near relative of the Modifiable Areal Unit Problem (MAUP, 15–17). The MAUP and the problem of choosing appropriate aggregating units has a long history in the literature predating even Robinson's ecological fallacy by at least 15 years. MAUP, was first recognized in a 1934 study by Gehlke and Biehl (18), and has since been well-documented in a range of studies (19–23). The term MAUP encompasses two closely related issues: zoning and scale (16). Zoning refers to dividing an area into the same number of units in different sizes, shapes, and configurations. Scale refers to changes in the number of units in an area. Changes of both types can influence observed relationships and require analysts to think carefully about the areal units for representing context. Fotheringham and Wong, for example, tested Census block groups and tracts to represent different zoning possibilities and different contextual scales in Buffalo, New York. They conclude that in the case of multivariate analysis "[c]alibration results from one set of areal units are highly suspect and should not be relied upon to draw any substantive conclusions about the underlying relationships being examined" (16).

Zoning and scale produce different problems with respect to how they influence contextual variables. With respect to scale, units that are too large may mis-specify context by missing important within-unit variation or by combining regions of two or more types together to produce an average value that describes neither component region (24, 25). Conversely, units that are too small may mis-specify context by including local variation that is not meaningfully experienced by residents of a particular context—in some circumstances, the effect may even be entirely missed by the analysis. Small aggregating units have a related problem of being highly sensitive to where boundaries are drawn; the typically smaller populations mean that the inclusion or exclusion of just a few individuals can alter the reported context significantly. In the absence of an administratively meaningful unit of observation, analysts are faced with a difficult task of finding units that are small enough to pick up variation in context, but large enough to avoid meaningless variation. Researchers, who rarely have control over where and how aggregating unit boundaries are drawn, must simply hope that the available data do not break up or combine meaningful contexts leaving them with variables that do not match the experience they hope to model. When contexts are inevitably somewhat mis-specified then researchers also must hope that whatever variance arises from the definition of boundaries is random rather than geographically systematic (26).

MAUP is a particular problem if a relevant variable varies across geographic areas. In a broadly homogenous population, choices related to scale or zoning may have little effect on contextual variables, but in the presence of a heterogeneous population these choices may have substantial effects. For example, when considering percent foreign born as a contextual variable in the study of individual views on immigration (27), the choice to use zip codes or

census tracts as an aggregating unit may have little effect on the value assigned to context in the absence of a substantial foreign born population. When immigrant populations are large and/or clustered in space, however, the potential for error is considerable. Even the scale at which similar population sorting processes operate can vary, as Johnston and co-authors have shown with respect to populations with different ancestry in London and Sydney (28, 29). If sorting occurs at different scales in different parts of a study region, then variation in the quality of the specification of context will not only impact research results, but also introduce systematic geographic bias as well.

Ultimately, MAUP makes assigning an appropriate context impossible in the absence of some administratively meaningful boundary, but it is possible to estimate the magnitude of effects and their potential impact on results. From Robinson, the idea of a *contextual fallacy* implies that individuals within the same aggregating unit may experience different contexts. A value capturing the degree to which individual context varies within a unit conveys the degree to which the zoning and scale choices underlying a given set of observations are likely to affect the context that is ultimately assigned to individuals. Where there is little variation across individuals within an observation, MAUP is probably not an important issue; where individual context varies, MAUP may be quite significant.

The contribution of this paper is to quantify the degree to which individual experience varies within common aggregating units often used to represent individual context in the United States. We propose a simple measure; the *standard deviation of individual context* (SDIC hereafter), that compares the local environment of individuals to the context assigned to them. It thereby captures the degree to which a given contextual value is a good fit for the individuals to which it is assigned. While the measure itself is straightforward, most scholars lack access to individual-level data from the U.S. Census. We compute an SDIC, therefore, for three Census geographies (the block, block group, and tract) and for three common contextual variables (percent Black, percent Hispanic, and Entropy). Strict Census rules to insure individual privacy protection appropriately prohibit release of these small-area values, but allow county-level averages that can be used to test the degree to which the contextual fallacy may condition outcomes.

We find sizeable differences in the level of variation between block and block group measures, but little change when moving between block groups and tracts. We find systematic spatial variation across U.S. regions with the highest variation present in places with high percentages of the population being treated as context (e.g. Hispanics in the Southwest and Blacks in the South). The uneven geographic variation in SDIC indicates that these contextual variables have the potential to introduce an uneven spatial process into results when the contextual fallacy is not considered.

Data

The results presented here are derived from nearly all of the 308,745,538 individual responses to the 2010 U.S. Decennial Census with each response matched to latitude and longitude coordinates assigned to place of residence. We code each individual into one of five ethno-racial categories: Non-Hispanic White, Non-Hispanic Black, Non-Hispanic Asian, Hispanic, and Non-Hispanic Other (White, Black, Asian, Hispanic, and Other hereafter). Other contains Non-Hispanic individuals self-identifying as Native American, Pacific Islander or “other” race as well as individuals identifying as two or more races. Each individual is further associated with the Census block, block group, and tract to which they were assigned for the publicly available 2010 tabulations.

The data were made available within the secure Census Federal Statistical Research Data Center (RDC) environment where their use is constrained to protect the privacy of individual responses. Consequently, we cannot release publicly the exact numbers of individuals used in the analysis, although we can confirm that the sample used exceeds 95% of the population. We have omitted the small number of individuals for whom no coordinates are reported, and we follow common practice in the literature on residential sorting by omitting individuals living in group quarters. The results presented below are aggregated in a variety of ways to protect individual privacy. The complete tract, block group, and block level results are only available to researchers within the protected context of the RDC, although the Census Bureau has determined no disclosure risk from a small subset of tract measures (with tracts having fewer than 100 residents omitted entirely) and for the data aggregated to the county level. The county-level SDIC scores represent the population-weighted average SDIC for all units within a county.

Results

For this analysis we create egocentric neighborhoods centered on each individual in the US (30) as a measure of *individual context* that can be used as a point of comparison to determine the variation from publicly available tabulations for Census geographies like the block or tract. We refer to these publicly available values hereafter as the *unit context* to generalize our comparison to any Census aggregating unit. Specifically, for any individual i , the *individual context* for variable a is defined with reference to the k nearest neighbors of i in Euclidean space. Thus, a_{ik} for percent Black would refer to the share of i 's k neighbors who self-identified as Black in the Census. The individual context is defined relative to a specific unit context so that k equals the number of individuals tabulated in the unit context. Consequently, a_{ik} defining the individual context at the block level for an individual living in a block with 50 residents would be based on the proportion of individual i 's 50 nearest neighbors whether or not they resided in the same block as i .

Figure 1 uses 250 points assigned to be either Black or White to illustrate the concepts of *unit context* and *individual context*. In panel A, all points are randomly assigned a race and

location. Fifty-four individuals highlighted in green are located inside the polygon, thirty of those individuals are Black yielding a unit context for percent Black in the polygon $a_c = 30/54 = 56\%$. Note that this unit context measure applies to all of the points highlighted in green regardless of where they fall within the polygon. In panel B the fifty-four nearest neighbors of the individual i denoted by ∇ are highlighted in blue. Twenty-eight of these individuals are Black giving an individual context for percent Black $a_{ic} = 28/54 = 52\%$. Panel C shows interpolated values for a_{ic} for the random data used in panels A and B. Panel D extends the example by introducing a population that is not randomly placed but structured so that black individuals are slightly more likely to appear to the left of the study region with Whites slightly more likely to appear more to the right. In Panel C a_{ic} varies by $\pm 10\%$ around a_c within the polygon. In Panel D, with just a small amount of spatial sorting introduced, a_{ic} varies by $\pm 25\%$ depending on where in the polygon it is measured.

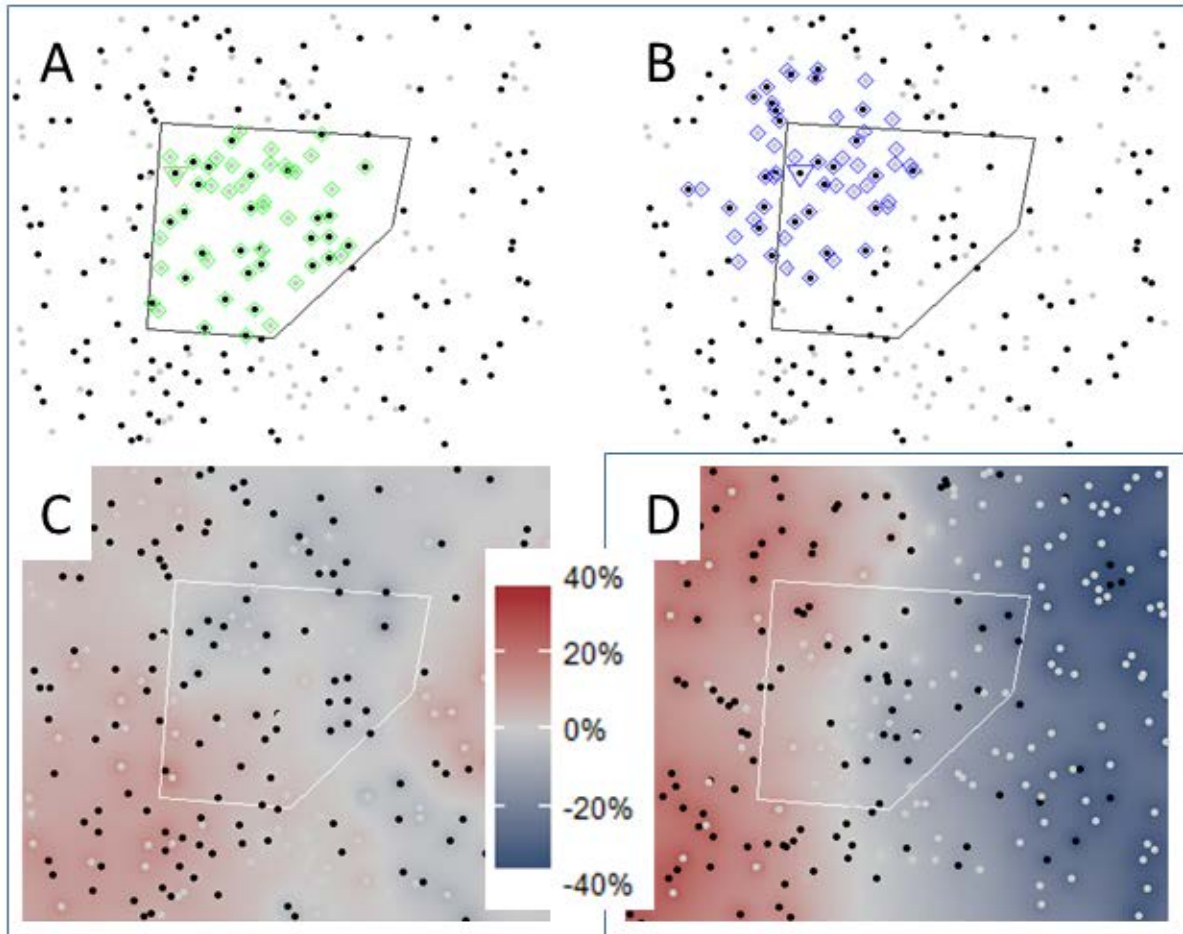


Figure 1: Worked example of the variation in individual context for a single aggregating unit with randomly assigned individuals. (A) individuals who compose the unit context for the polygon, (B) individuals who constitute the individual context for the individual denoted by ∇ , (C) variation in individual context using same points as in panels A and B, (D) variation in individual context when points are assigned a racial category with spatial clustering.

Table 1 provides a summary for the entire U.S. of the extent of variation in individual experience for three key measures at three important Census geographies. The table contains some very specific answers to basic questions about the magnitude of the contextual fallacy as a problem with respect to *zoning*. When conducting tract-level analyses, on average only 5.7% and 4.2% of individuals vary by more than 10% from the value for Percent Black and Percent Hispanic assigned to their tract, and less than 1% of the population vary from the unit value by more than 25% in both cases. The contextual fallacy is a greater concern in analyses of Entropy, i.e., population diversity, where within-tract variation of 10 points or more typically occurs for 11% of the population although only 1% of the population live in a context that is 25 points more or less diverse than the contextual value to which they would normally be assigned.

Table 1 also suggests some striking properties with respect to variation and *scale*. The move from tract to block group represents a significant shrinking in the number of neighbors included in a context, but is associated with almost no change in the percentages shown in Table 1. In contrast, the move to a block-level context, which constitutes another reduction in the number of neighbors, radically increases the number of individuals who experience a context dramatically different from the one to which they are assigned in the Census. From this table it appears that the individual context may vary too much at the block level to make blocks a good unit of observation for these contextual variables at this scale. In contrast, the small difference between tracts and block groups indicates on the one hand that tracts are probably not over-generalizing context by grouping dissimilar contexts together; on the other hand, block groups will give a finer geographic resolution without succumbing to the excessive variation plaguing blocks. The fact that over 9% of the population varies by at least 25 points from their block level entropy score seems to indicate that this measure should not be used as a contextual variable with block-level observations.

Table 1: Variation in individual experience with respect to reported context variable values

	Percent of individuals who differ from the unit measure by at least					
	10%			25%		
	Tract	Block Group	Block	Tract	Block Group	Block
Percent Black	5.7	5.1	10	0.8	0.8	2.1
Percent Hispanic	4.2	4.0	14.5	0.4	0.4	2.9
Entropy	11.0	10.9	36.0	1.1	1.0	9.4

Percentages based on complete sample of U.S. 2010 Decennial Census respondents living in administrative units containing at least 10 individuals. Number of units suppressed to meet Census privacy requirements

Table 1 is informative in describing the magnitude and structure of the contextual fallacy, but to be useful to researchers it is important to generate a measure that will permit them to test whether or not the contextual fallacy has the potential to alter their substantive findings. To this end we further propose the *standard deviation of individual context (SDIC)*, represented here by sigma (σ). For our analysis we calculate sigma for every block, block group, and tract in the 2010 Census. Sigma is defined for a specific unit context c (a block, block group, or tract) and contextual variable a (percent Black, percent Hispanic, and Entropy²) and captures the degree to which individual context varies from unit context for all of the individuals who

² We follow Reardon and Firebaugh (32) in defining Entropy for our five ethno-racial categories as:

$E = \left(\sum_{m=1}^M \pi_m \ln(1/\pi_m) \right) / \ln(M)$ where m denotes one of our five ethno-racial categories, π_m is the proportion of the total population in category m , and \ln refers to the natural log with $\ln(1/\pi_m)$ treated as 0 when π_m equals 0. By scaling our entropy values by $\ln(M)$, we constrain the range of E from 0 to 1. When E equals 1, diversity is maximized: all ethno-racial groups are identical in size. A zero value, on the other hand, indicates the absence of diversity (homogeneity) such that all residents belong to the same group.

reside in a specific unit. For any contextual variable a_c for census unit c composed of k individuals i , σ_{ac} is defined as:

$$\sigma_{ac} = \sqrt{\frac{\sum_{i \in c} (a_{ik} - a_c)^2}{k}}$$

Figure 2 illustrates how SDIC varies at the tract level for the region in and around the City of Seattle.³ In some areas, census tract boundaries create few issues with respect to the contextual fallacy. Specifically, individuals in the predominantly white tracts of North Seattle are well matched to their census tract contexts with respect to Percent Black and Percent Hispanic. This is also true of similar areas just to the east of the City of Seattle on Mercer Island and along the eastern shore of Lake Washington. In other parts of the region, however, SDIC increases substantially: exceeding 15% for some tracts. Because SDIC is a measure of standard deviation, a value of 15% indicates that roughly a third of the population has a mis-specified context of 15% or more. Values of SDIC with respect to entropy (Figure 2C) vary more but do so in the Seattle case with less evidence of clustering in space.

³ For privacy protection Census will not release tract-level measures of sigma for tracts with fewer than 100 people. Seattle was chosen as an example because it is well known to the researchers and could therefore be interpreted with greater accuracy, it has no special properties with respect to sigma in other respects.

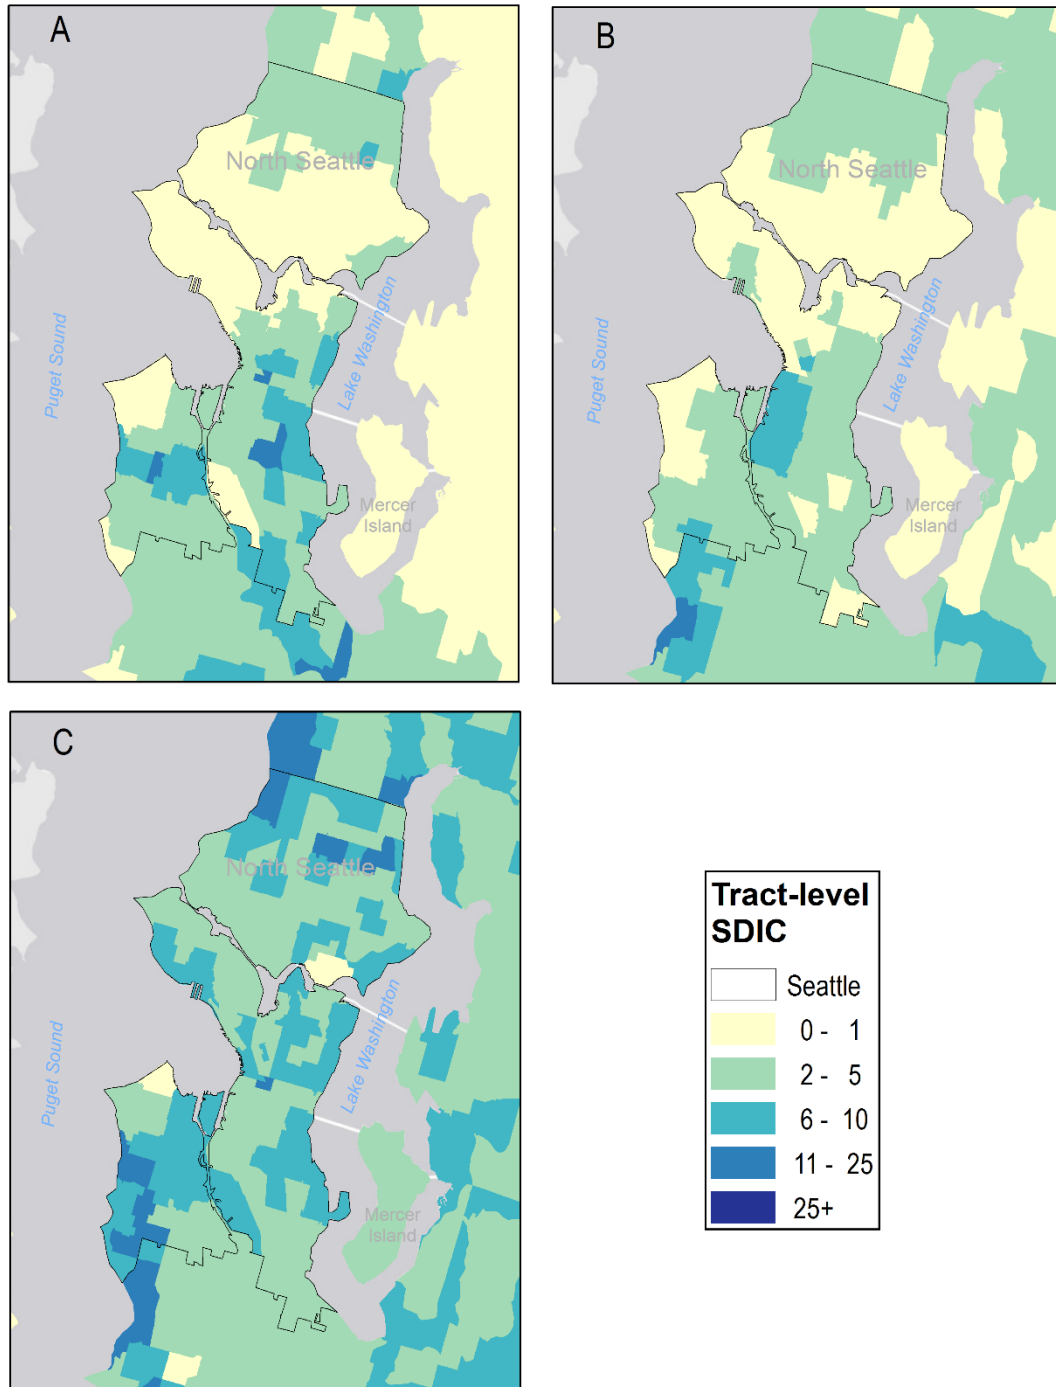


Figure 2 Tract-level SDIC for three contextual variables: A) Percent Black, B) Percent Hispanic, and C) Entropy. City of Seattle and surrounding area. Service layer credits: ESRI, HERE, DeLorme, MapmyIndia, © OpenStreetMap contributors, and the GIS user community

Turning to county averages of SDIC for the entire nation, which are available at [<URL to be added>](#), Figure 3 maps the county averages for tract-level SDIC across our three contextual variables. Figures 3A and 3B notably indicate that, for large portions of the U.S., average tract-

level SDIC is not very high. The maps also indicate that the average values presented in Table 1 mask the comparatively high SDIC in places with significant populations of Blacks and Hispanics. The unevenness in variation from unit context indicates that contextual effects would be measured with uneven quality for national-level analysis and introduces systematic geographic variation in measure quality. In multivariate analysis, in which variables for both percent Hispanic and percent Black might be employed, the regional variation in measure quality could have complicated interaction effects in places like North Carolina where moderate values of SDIC are present on both measures. In contrast to Panels 3A and 3B, Panel 3C shows that there is no geographic pattern for variation in Entropy (similar to Seattle), but that the average SDIC is considerably higher for Entropy than it is for percent Black and percent Hispanic. This suggests that tract level Entropy measures may not introduce systematic spatial problems into analyses, but that moderate variance is present in most of the country.

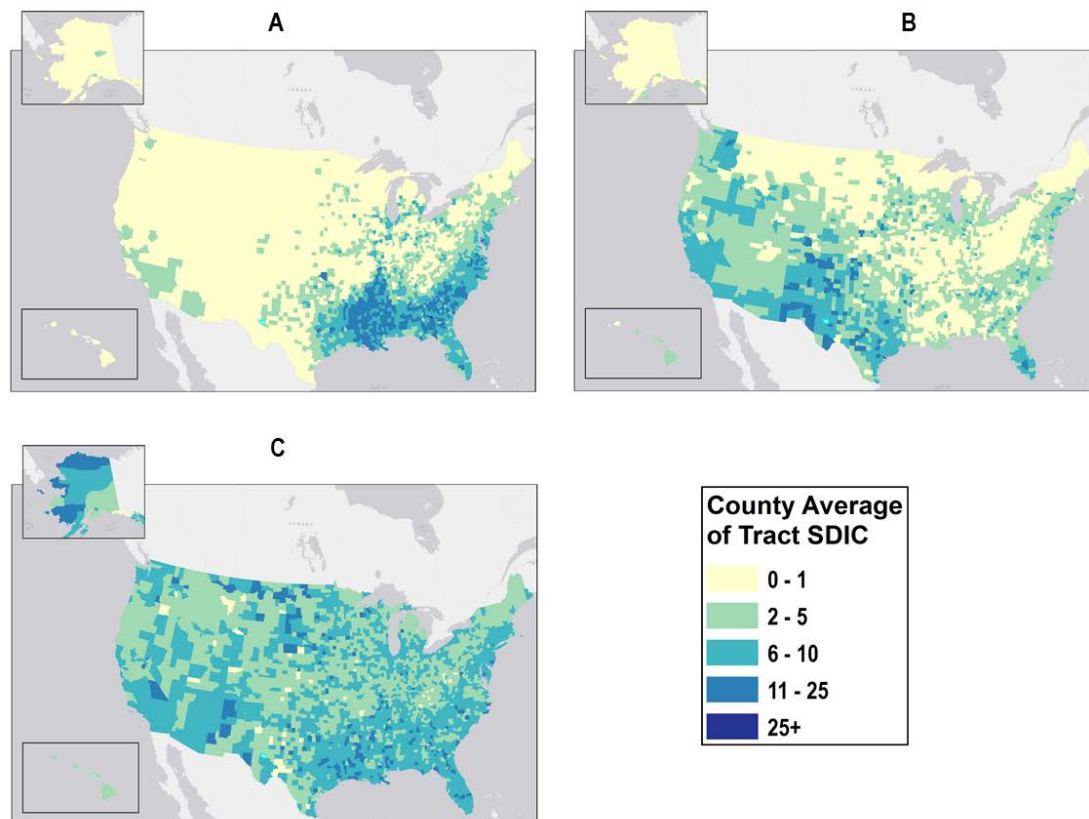


Figure 3: County average tract-level sigma for three contextual variables. (A) Percent Black, (B) Percent Hispanic, and (C) Entropy. Service layer credits: ESRI, HERE, DeLorme, MapmyIndia, © OpenStreetMap contributors, and the GIS user community

Using SDIC

The SDIC can help determine whether the contextual fallacy plays a significant role in shaping outcomes in a given analysis. The SDIC is an indicator of potential problems related to the contextual fallacy, but its significance with respect to a specific research question will

depend on a number of factors including the study area, strength of effect, number of variables and structure of the analysis. The SDIC will not be useful for some complicated modeling strategies, but its application to linear regression is relatively straightforward. Because it represents a standard deviation, the SDIC can be used with linear regression in conjunction with methods quite similar to those used for multiple imputation of missing data (31). Using the published *unit* value for a given block, block group, or tract as the mean and the scale-appropriate SDIC as the standard deviation it is possible to generate multiple sets of synthetic data. These synthetic data can then be analyzed and the results of these multiple analyses pooled to determine the effect of the contextual fallacy on results. If more than one of the contextual variables defined here is included in a model, then tests will have to be run separately to avoid issues associated with sampling in the context of covariance.

The block, block group, and tract-level SDIC measures are available for every geography in the United States within the secure environment of the RDC, conditional of course on Census approval to use those data. The values mapped in Figure 3 (and their equivalents for blocks and block groups) are available from <URL to be added> and represent the average SDIC across all subunits in a county. While not as exact as the unit-level measures of SDIC, these values can help researchers assess whether the contextual fallacy has the capacity to alter their substantive findings. To facilitate the implementation of this method, a worked example using the open-source software program R is distributed in conjunction with the SDIC values at the above web site.

Conclusion

Context matters for a host of social science research questions, but poses issues of measurement error for scholars. The definition of boundaries for relevant populations and their sorting in space can profoundly shape the context as experienced by any given individual. A remedy for this problem is unattainable because no delineation will perfectly capture the context for all individuals. However, scholars can estimate the degree to which contextual variables may be mis-specified in their analysis to assess potential effects on their substantive results. Based on individual responses to the 2010 U.S. Decennial Census, the standard deviation of individual context provides the means to undertake these tests. Its uneven distribution in space for both percent Black and percent Hispanic suggests that it may be useful in detecting systematic geographic error in findings. Further applications of the SDIC will be necessary to determine thresholds at which variation in individual context is high enough to impact analytic results. Moreover, differences between tract, block group, and block-level measures of SDIC could be employed as a means to better understand the effect of using these geographic units to aggregate populations into a measure of context.

References

1. Robinson WS (1950) Ecological correlations and the behaviors of individuals. *Am Sociol Rev* 15:351–357.
2. Firebaugh G (1978) A Rule for Inferring Individual-Level Relationships from Aggregate Data. *Am Sociol Rev* 43(4):557.
3. Subramanian S, Jones K, Kaddour A, Krieger N (2009) Revisiting Robinson: the perils of individualistic and ecologic fallacy. *Int J*.
4. Ludwig J, et al. (2013) Long-term neighborhood effects on low-income families: Evidence from moving to opportunity. *American Economic Review* doi:10.1257/aer.103.3.226.
5. Sharkey P, Faber JW (2014) Where, When, Why, and For Whom Do Residential Contexts Matter? Moving Away from the Dichotomous Understanding of Neighborhood Effects. *Annu Rev Sociol* 40(1):559–579.
6. Goering J (2013) Neighborhood Effects and Public Policy. *City Community* 12(1):13–20.
7. Ludwig J, et al. (2008) What can we learn about neighborhood effects from the Moving to Opportunity experiment. *Am J Sociol* 114(1):144–188.
8. Clampet-Lundquist S, Massey D (2008) Neighborhood Effects on Economic Self-Sufficiency: A Reconsideration of the Moving to Opportunity Experiment1. *Am J Sociol*.
9. Sampson RJ, Morenoff JD, Gannon-Rowley T (2002) Assessing "Neighborhood Effects:: Social Processes and New Directions in Research. *Annu Rev Sociol* 28:443–478.
10. Galster GC (2008) Quantifying the Effect of Neighbourhood on Individuals: Challenges, Alternative Approaches, and Promising Directions. *Schmollers Jahrb* 128(1):7–48.
11. Colabianchi N, et al. (2007) Towards an understanding of salient neighborhood boundaries: adolescent reports of an easy walking distance and convenient driving distance. *Int J Behav Nutr Phys Act* 4(1):66.
12. Campbell E, Henly JR, Elliott DS, Irwin K (2009) SUBJECTIVE CONSTRUCTIONS OF NEIGHBORHOOD BOUNDARIES: LESSONS FROM A QUALITATIVE STUDY OF FOUR NEIGHBORHOODS. *J Urban Aff* 31(4):461–490.
13. Plane DA (1981) The geography of urban commuting fields: Some empirical evidence from new England. *Prof Geogr* 33(2):182–188.
14. Fowler CS, Rhubart DC, Jensen L (2016) Reassessing and revising Commuting Zones for 2010: History, assessment, and updates for U.S. "labor-sheds" 1990-2010. *Popul Res Policy Rev* 35(2):263–286.
15. Openshaw S (1977) A geographical solution to scale and aggregation problems in region-building, partitioning and spatial modelling. *Trans Inst Br Geogr* 2(4):459–472.
16. Fotheringham A, Wong D (1991) The modifiable areal unit problem in multivariate statistical analysis. *Environ Plan A* 23(7):1025–1044.
17. Gryparis A, Paciorek CJ, Zeka A, Schwartz J, Coull BA (2009) Measurement error caused

- by spatial misalignment in environmental epidemiology. *Biostatistics* 10(2):258–274.
18. Gehlke CE, Biehl K (1934) Certain effects of grouping upon the size of the correlation coefficient in census tract material. *J Am Stat Assoc* 29(185):169–170.
 19. Larson R (1986) The Invalidity of Modifiable Areal Unit Randomization. *Prof Geogr* 38(4):369–374.
 20. Openshaw S, Taylor PJ (1979) A million or so correlation coefficients: three experiments on the modifiable areal unit problem. *Statistical Applications in the Spatial Sciences*, pp 127–144.
 21. Openshaw S (1984) Ecological fallacies and the analysis of areal census data. *Environ Plan A* 16(1):17–31.
 22. Openshaw S, Taylor PJ (1981) The modifiable areal unit problem. *Quantitative Geography: A British View*, eds Wrigley N, Bennett RJ (Routledge, London), pp 60–69.
 23. Wong D (2004) The modifiable areal unit problem (MAUP). *Worldminds Geogr Perspect* 100.
 24. Kropp P, Schwengler B (2016) Three-Step Method for Delineating Functional Labour Market Regions. *Reg Stud* 50(3):429–445.
 25. Krupka DJ (2007) Are Big Cities More Segregated? Neighbourhood Scale and the Measurement of Segregation. *Urban Stud* 44(1):187–197.
 26. Arbia G (2006) *Spatial Econometrics* (Springer-Verlag, Berlin/Heidelberg) doi:10.1007/3-540-32305-8.
 27. Citrin J, Green DP, Muste C, Wong C (1997) Public Opinion Toward Immigration Reform: The Role of Economic Motivations. *J Polit* 59(3):858.
 28. Johnston R, Jones K, Manley D, Owen D (2016) Macro-scale stability with micro-scale diversity: modelling changing ethnic minority residential segregation – London 2001–2011. *Trans Inst Br Geogr* 41(4):389–402.
 29. Johnston R, Forrest J, Jones K, Manley D (2016) The scale of segregation: ancestral groups in Sydney, 2011. *Urban Geogr* 37(7):985–1008.
 30. Spielman SE, Logan JR (2013) Using High-Resolution Population Data to Identify Neighborhoods and Establish Their Boundaries. *Ann Assoc Am Geogr* 103(1):67–84.
 31. Little R, Rubin D (2014) *Statistical analysis with missing data*.
 32. Reardon SF, Firebaugh G (2002) Measures of Multigroup Segregation. *Sociol Methodol* 32(1):33–67.

Figure Legends

Figure 1: Worked example of the variation in individual context for a single aggregating unit with randomly assigned individuals. (A) individuals who compose the unit context for the polygon, (B) individuals who constitute the individual context for the individual denoted by ∇ ,

(C) variation in individual context using same points as in panels A and B, (D) variation in individual context when points are assigned a racial category with spatial clustering.

Figure 2 Tract-level SDIC for three contextual variables: A) Percent Black, B) Percent Hispanic, and C) Entropy. City of Seattle and surrounding area. Service layer credits: ESRI, HERE, DeLorme, MapmyIndia, © OpenStreetMap contributors, and the GIS user community

Figure 3: County average tract-level sigma for three contextual variables. (A) Percent Black, (B) Percent Hispanic, and (C) Entropy. Service layer credits: ESRI, HERE, DeLorme, MapmyIndia, © OpenStreetMap contributors, and the GIS user community